

Two-Dimensional Pursuit-Evasion in a Compact Domain with Piecewise Analytic Boundary

Andrew Beveridge* and Yiqing Cai†

Abstract

In a pursuit-evasion game, a team of pursuers attempt to capture an evader. The players alternate turns, move with equal speed, and have full information about the state of the game. We consider the most restrictive capture condition: a pursuer must become colocated with the evader to win the game. We prove two general results about pursuit-evasion games in topological spaces. First, we show that one pursuer has a winning strategy in any $\text{CAT}(0)$ space under this restrictive capture criterion. This complements a result of Alexander, Bishop and Ghrist, who provide a winning strategy for a game with positive capture radius. Second, we consider the game played in a compact domain in Euclidean two-space with piecewise analytic boundary and arbitrary Euler characteristic. We show that three pursuers always have a winning strategy by extending recent work of Bhaduria, Klein, Isler and Suri from polygonal environments to our more general setting.

1 Introduction

A pursuit-evasion game in a domain \mathcal{D} is played by a team of pursuers p_1, p_2, \dots, p_k and an evader e . The pursuers win if some p_i becomes colocated with the evader after a finite number of turns, meaning that the distance $d(p_i, e) = 0$. When this occurs, we say that p_i *captures* e . We consider the discrete time version of the game, which proceeds in turns. Initially, the pursuers choose their positions $p_1^0, p_2^0, \dots, p_k^0$, and then the evader chooses his initial position e^0 . In turn $t \geq 1$, each pursuer p_i moves from her current position p_i^{t-1} to a point $p_i^t \in B(p_i^{t-1}, 1) = \{x \in \mathcal{D} \mid d(p_i^{t-1}, x) \leq 1\}$. If the evader has been captured, then the game ends with the pursuers victorious. Otherwise, the evader moves from e^{t-1} to a point $e^t \in B(e^{t-1}, 1)$. The evader wins if he remains uncaptured forever. We consider the full-information (full-visibility) game in which each player knows the environment and the locations of all the other players. Furthermore, the pursuers may coordinate their movements.

Turn-based pursuit games in simply connected domains have been well-characterized: one pursuer is sufficient to capture the evader. Winning pursuer strategies have been found for environments in \mathbb{R}^n [21, 16], and in simply connected polygons [14]. Taking a topological viewpoint and using the

*Department of Mathematics, Statistics and Computer Science, Macalester College, Saint Paul, MN 55105. abeverid@macalester.edu

†The Institute for Mathematics and Its Applications, University of Minnesota, Minneapolis, MN 55455, yiqingcai@ima.umn.edu

weaker capture criterion $d(p, e) < \epsilon$ for some constant $\epsilon > 0$, Alexander, Bishop and Ghrist [3] proved that a single pursuer can capture an evader in any compact $\text{CAT}(0)$ by heading directly towards the evader at maximum speed. We provide an alternate strategy for a compact $\text{CAT}(0)$ space that achieves $d(p, e) = 0$ in a finite number of turns. Our winning pursuer strategy is the topological version of *lion's strategy*, which has been used successfully in \mathbb{E}^n [21] and in simple polygons [14]. We defer the description of this strategy to Section 2.

Theorem 1.1. *A pursuer p using lion's strategy in a compact $\text{CAT}(0)$ space \mathcal{D} , captures the evader e by achieving $d(p, e) = 0$ after at most $\text{diam}(\mathcal{D})^2$ turns.*

Theorem 1.1 implies that a single pursuer can become colocated with an evader in a simply connected, compact domain $\mathcal{D} \subset \mathbb{E}^2$. But it is easy to construct compact domains that are evader win: removing one large open set from the middle of a simply connected domain tips the game in the evader's favor. Indeed, the evader can keep this large obstruction between himself and the pursuer, indefinitely. Such an open set is called an *obstacle* or *hole* in the environment. It is not hard to show that adding a second pursuer to this two-dimensional domain gives the game back to the pursuers. Adding multiple obstacles creates a distinct topology, and it is natural to wonder how many pursuers are needed to capture an evader in such an environment. The analogous question has been resolved for pursuit-evasion games in certain two-dimensional environments. Aigner and Fromme [1] proved that three pursuers are sufficient for pursuit-evasion on a planar graph. Bhadauria, Klein, Isler and Suri [8] showed that the analogous result holds in a two-dimensional polygonal environment with polygonal holes. We generalize this three-pursuer result to a broader class of topological spaces.

Our pursuit game takes place in a compact and path-connected domain $\mathcal{D} \subset \mathbb{E}^2$. The set \mathcal{D} contains a finite set of disjoint open obstacles $\mathcal{O} = \{O_1, O_2, \dots, O_k\}$. The boundary of the domain is $\partial\mathcal{D} = \{\partial O_0, \partial O_1, \partial O_2, \dots, \partial O_k\}$ where we define ∂O_0 to be the outer boundary of \mathcal{D} , for convenience. We place two conditions on the boundary. First, $\partial\mathcal{D}$ can be decomposed into a finite number of analytic curves $\gamma(t) = (x(t), y(t))$ for $0 \leq t \leq 1$, where each of $x(t), y(t)$ can locally be expanded as convergent power series. Second, we require that $\partial\mathcal{D}$ is a 1-manifold: for any $x \in \partial\mathcal{D}$, there exists an $\epsilon > 0$, such that $B(x, \epsilon) \cap \partial\mathcal{D}$ is homeomorphic to \mathbb{R}^1 . In other words, we forbid self-intersections. For brevity, we say that a domain \mathcal{D} satisfying these properties is *piecewise analytic*. We list three consequences of these conditions. First, the number of singular points on the boundary is finite. Second, the absolute value of the curvature at the nonsingular points of $\partial\mathcal{D}$ is bounded above by some constant $\kappa_{\max} > 0$. Third, there is a minimum separation $d_{\min} > 0$ between boundary components: $d(O_i, O_j) > d_{\min}$ for all $0 \leq i < j \leq k$. During the game, the pursuers will guard a sequence of geodesics; crucially, we will see in Section 4 that each of these geodesics is also piecewise analytic. This brings us to our main result.

Theorem 1.2. *Three pursuers can capture an evader in a compact domain in \mathbb{E}^2 with piecewise analytic boundary. The number of turns required to capture the evader for a domain with k obstacles is*

$$O(2k \cdot \text{diam}(\mathcal{D}) + \text{diam}(\mathcal{D})^2).$$

At a high level, our winning three-pursuer strategy builds directly on those found in [1, 8], and we are indebted to those previous papers. However, our geometric and topological approach is entirely new. In particular, our arguments are grounded in a careful investigation of the convexity, curvature and homotopy classes of geodesic curves in our domain. And of course, Theorem 1.2 significantly extends the class of known three-pursuer-win domains.

1.1 Related Work

Pursuit-evasion games have enjoyed a long research history. In the 1930s, Rado posed the Lion and Man game in which a lion hunts the man in a circular arena. The players move with equal speeds, and the lion wins if it achieves colocation. At first blush, it seems that lion should be able to capture man, regardless of the man’s evasive strategy. However, Besicovitch showed that when the game is played in continuous time, the man can follow a spiraling path so that lion can get arbitrarily close, but cannot achieve colocation [17]. However, when lion and man move in discrete time steps, our intuition prevails: lion does have a winning strategy [21].

The interdisciplinary literature on pursuit-evasion games spans a range of settings and variations. Pursuit games have been studied in many environments, including graphs, in polygonal environments and in topological spaces. Researchers have considered motion constraints such as speed differentials between the players, constraints on acceleration, and energy budgets. As for sensing models, the players may have full information about the positions of other players, or they may have incomplete or imperfect information. Typically, the capture condition requires achieving colocation, a proximity threshold, or sensory visibility (such as a non-obstructed view of the evader). For an overview of pursuit-evasion on graphs, see the monograph [9]. The papers [11] and [16] provide a nice introduction to pursuit in the polygonal setting.

The past decade has witnessed a renaissance of pursuit-evasion results in multiple disciplines. Prominent research efforts come from the robotics community, where pursuit-evasion in polygonal environments is a productive setting for exploring autonomous agents. Pursuit-evasion has also thrived in the graph theory community, where it is known as the game of Cops and Robbers. More recently, researchers have started exploring pursuit-evasion games in topological spaces. This is a natural evolution for the study of pursuit-evasion games. Indeed, determining the number of pursuers required to capture an evader in a given environment becomes a question about its topology since the various loops and holes of the environment provide escape routes for the evader.

The classic paper of Aigner and Fromme [1] initiated the study of multiple pursuers versus a single evader on a graph. In this turn-based game, agents can move to adjacent vertices, and the cops win if one of them becomes co-located with the robber. This paper introduced the *cop number* of a graph, which is the minimum number of pursuers (cops) needed to catch the evader (robber). Aigner and

Fromme proved that the cop number of any planar graph is at most 3. This bound is tight, as the dodecahedron graph requires three cops. At a high level, their winning pursuer strategy proceeds as follows. Two cops guard distinct (u, v) -paths where u, v are vertices of the graph G . This restricts the pursuer movement to a subgraph of G . The third pursuer then guards another (u, v) -path, chosen so that (1) the robber’s movement is further restricted, and (2) one of the other cops no longer needs to guard its path. This frees up that cop to continue the pursuit. This process repeats until the evader is caught.

More recently, an analogous result was proven by Bhadauria, Klein, Isler and Suri [8] for pursuit-evasion games in a two-dimensional polygonal environment with polygonal holes. In this turn-based game, an agent can move to any point within unit distance of its current location. Like Aigner and Fromme, they use colocation as their capture criterion. Bhadauria et al. prove that three pursuers are sufficient for pursuit-evasion in this setting, and that this bound is tight. The pursuer strategy is inspired by the Aigner and Fromme strategy for planar graphs: two pursuers guard paths that confine the evader while the third pursuer takes control of another path that further restricts the evader’s movement. Of course, the details of the pursuit and the technical proofs are quite different from the graph case. Their proofs make heavy use of the polygonal nature of the environment, both to find the paths to guard and to guarantee that their pursuit finishes in finite time.

Just as the proofs of Bhadauria et al. were inspired by Aigner and Fromme, our proof of Theorem 1.2 is inspired by those for the polygonal environment. Bhadauria et al. actually give two different winning strategies for three pursuers. At a high level, these strategies progress in the same way, but the tactics for choosing paths and how to guard them are different. Herein, we adapt their *shortest path strategy* to our more general setting. The topological environment introduces a distinctive set of challenges to overcome. In particular, we do not have a finite set of polygonal vertices to use as a backbone for our guarded paths. Instead, we rely on homotopy classes to differentiate between paths to guard. Looking beyond the high-level structure of our pursuer strategy, the arguments (and their technical details) in this paper are wholly distinct from those found in [8], and our result applies to a much broader class of environments.

Finally, we note that our result follows in the footsteps of other recent explorations of pursuit-evasion games in general geometric and topological domains. Pursuit-evasion games in such spaces have applications in robotics, where agents must navigate and coordinate in high dimensional configuration spaces. Alexander, Bishop and Ghrist helped to pioneer this subject, studying pursuit-evasion games with the capture condition $d(p, e) < \epsilon$ for some constant $\epsilon > 0$ (rather than colocation). In [3], Alexander, Bishop and Ghrist prove that a single pursuer can capture an evader in any compact $\text{CAT}(0)$ space: the *simple pursuit* strategy of heading directly towards the evader is a winning strategy. In [5], these authors explore the simple pursuit strategy in unbounded $\text{CAT}(K)$ spaces with positive curvature $K > 0$, developing connections between evader-win environments and the total curvature of the pursuer’s trajectory. Finally, in [4], they consider pursuit games in unbounded Euclidean domains using

multiple pursuers. They provide conditions on the initial configuration of the players that guarantee capture, generalizing (and amending) results of Sgall [21] and Kopparty and Ravishankar [16].

1.2 Preliminaries

We introduce some notation and review some concepts and results from algebraic topology [13]. We then prove three lemmas about convex paths in two-dimensional compact regions with piecewise analytic boundary.

A topological space is a set X along with a collection of subsets of X , called *open sets* that satisfy a sequence of axioms [7]. A map $f : X \rightarrow Y$ between two spaces is *continuous* when the inverse image of every open set in Y is open in X . A *path* $\Pi : [0, 1] \rightarrow \mathcal{D}$ is a continuous map from interval $[0, 1]$ to \mathcal{D} , with initial point $\Pi(0)$ and terminal point $\Pi(1)$. The *length* $l(\Pi)$ of this path is its arc length in Euclidean space \mathbb{E}^2 . A path Π is a *loop* when $\Pi(0) = \Pi(1)$. A *simple path* has no self-intersections, meaning that Π is injective. By abusing notation, we write $x \in \Pi$ when $x = \Pi(t)$ for some $t \in [0, 1]$. For $x, y \in \Pi$, we use $\Pi(x, y)$ to denote the subpath connecting these points. The space X is *path-connected* if there exists a path between any pair of points $x, y \in X$.

A *homotopy* of paths is a family of maps $f_t : [0, 1] \rightarrow X$, $t \in [0, 1]$, such that the associated map $F : [0, 1] \times [0, 1] \rightarrow X$ given by $F(s, t) = f_t(s)$ is continuous, and the endpoints $f_t(0) = x_0$ and $f_t(1) = x_1$ are independent of t . The paths f_0 and f_1 are called *homotopic*. The relation of homotopy on paths with fixed endpoints is an equivalence relation and we use $[f]$ to denote the homotopy class of the curve f under this relation. The set $[f]$ of loops in X at the basepoint x_0 forms a group under path composition, called the *fundamental group* of X at the basepoint x_0 . The space X is *simply connected* when it is path-connected and its fundamental group is trivial. For example, a subspace X of \mathbb{E}^2 is simply connected if and only if it has the same homotopy type as a 2-disc.

We now turn to some geometric properties of paths in \mathbb{E}^2 . The distance $d(x, y)$ between points $x, y \in X$ is the length of a shortest (x, y) -path in X . When restricting ourselves to $R \subset X$, we use $d_R(x, y)$ to denote the distance between these points in the subdomain. We will frequently consider a subdomain R enclosed by two simple (u, v) -paths Π_1, Π_2 . We denote such a set as $R[\Pi_1, \Pi_2] \subset X$.

For a C^2 path $\gamma : [0, T] \rightarrow \mathbb{E}^2$, its curvature at $\gamma(t)$ is defined as $\kappa(t) = \pm \frac{\|\gamma'(t) \times \gamma''(t)\|}{\|\gamma'(t)\|^3}$, with the sign positive if the tangent turns counterclockwise, and negative if the tangent turns clockwise. The smoothness of a piecewise analytic curve $\gamma : [0, 1] \rightarrow \mathbb{E}^2$ ensures that its absolute curvature is bounded at its nonsingular points. If γ is piecewise C^2 and continuous, with $t_0 < t_1 < \dots < t_n$ as the preimages of the singular points, then its *total curvature* is

$$\kappa_{\text{total}}(\gamma) = \sum_{i=1}^n \int_{t_{i-1}}^{t_i} \kappa(t) dt + \sum_{i=1}^n \theta_i$$

where θ_i is the exterior angle at $\gamma(t_i)$, and $\theta_n = 0$ when $\gamma(t_0) \neq \gamma(t_n)$. This brings us to the celebrated Gauss-Bonnet Theorem which relates the total curvature of a closed curve with the Euler characteristic




Figure 1.1: (a) A piecewise convex (u, v) -path Π . (b) Shortcutting a non-convex (u, v) -path Π_1 .

of its enclosed region. In our setting, the Euler characteristic equals $1 - k$, where k is the number of obstacles in the region R .

Theorem 1.3. *[Gauss-Bonnet Theorem, cf. [12]] Given a compact region $R \subset \mathbb{E}^2$ with boundary ∂R , we have*

$$\kappa_{\text{total}}(\partial R) = 2\pi\chi(R),$$

where $\chi(R)$ is the Euler characteristic of R .

We use the Gauss-Bonnet Theorem to understand the effect of obstacles on shortest paths. In particular, we will consider pairs of paths Π_1, Π_2 with shared endpoints u, v . These paths will be piecewise analytic (so that they have a finite number of singular points). Our goal is to prove that if the shortest (u, v) -path is not unique, then each shortest path must touch an obstacle in the given region. We begin with a definition of convexity, which we define for the broader family of piecewise C^2 smooth curves; an example is shown in Figure 1.1(a).

Definition 1. *Let $\Pi : [0, 1] \rightarrow \mathbb{E}^2$ be a piecewise C^2 smooth curve in \mathbb{E}^2 . Then Π is convex when the following holds for any point $x \in \Pi \setminus \{u, v\}$:*

- (a) *If Π is C^2 smooth at x , then the curvature at x is nonpositive;*
- (b) *If Π is not C^2 smooth at x , then the tangent line at x turns clockwise by an angle $0 \leq \theta \leq \pi$.*

Lemma 1.4. *Let Π_1 and Π_2 be two piecewise analytic (u, v) -paths with $\Pi_1 \cap \Pi_2 = \{u, v\}$. If Π_1 is a shortest path in $R[\Pi_1, \Pi_2]$, and touches no obstacle inside $R[\Pi_1, \Pi_2]$, then Π_1 is convex in $R[\Pi_1, \Pi_2]$.*

Proof. We prove the lemma by contradiction. Suppose that there exists an $x \in \Pi \setminus \{u, v\}$ where the convexity of Π_1 in $R[\Pi_1, \Pi_2]$ is violated. Either (a) Π_1 is C^2 smooth at x , but the curvature at x is positive, or (b) Π_1 is not C^2 smooth at x , and the tangent line turns counterclockwise by an angle $0 < \theta < \pi$ at x , creating a non-convex corner. Let d_0 denote the minimum distance between Π_1 and any obstacle $O \in R[\Pi_1, \Pi_2]$ with $\Pi_1 \cap \partial O = \emptyset$.

Suppose that the curvature at x is positive, see Figure 1.1(b). There must be a C^2 subpath Π_x between y_1 and y_2 of Π_1 around x with positive curvature. Using the lower bound d_0 on the separation between Π_1 and any obstacles inside $R[\Pi_1, \Pi_2]$ and Π_2 , we may take y_1, y_2 to be close enough so that the line segment Λ connecting y_1 and y_2 lies inside $R[\Pi_1, \Pi_2]$ and does not encounter any obstacles. Replacing Π_x with Λ creates a path that is strictly shorter than Π_1 , contradicting the minimality of Π_1 .

Next suppose there is a non-convex corner at x . By an analogous argument to the previous case, we can create a short-cut Λ around x to make a shorter path than Π_1 , a contradiction. \square

Lemma 1.5. *Let Π_1, Π_2 be two (u, v) -paths with $\Pi_1 \cap \Pi_2 = \{u, v\}$. Suppose that Π_1 is a convex and piecewise analytic (u, v) -path in $R[\Pi_1, \Pi_2]$, and let Π_2 be a convex and piecewise analytic (v, u) -path in $R[\Pi_1, \Pi_2]$. Then Π_1 and Π_2 are both straight lines connecting u, v .*

Proof. Let $Q = Q[\Pi_1, \Pi_2]$, which is the closed region between Π_1 and Π_2 , pretending there being no obstacles. If we consider ∂Q , the concatenation of $\Pi_1(u, v)$ and $\Pi_2(v, u)$, it is in fact a loop bounding a simply connected region. By the Gauss-Bonnet Theorem 1.3, the total curvature along ∂Q equals $2\pi\chi(Q) \geq 2\pi$. We decompose the value $2\pi\chi(Q)$ as the sum of total curvature of Π_1 and Π_2 respectively, and the exterior angles at u and v . Because of convexity, both Π_1 and Π_2 have total curvature no greater than 0. As for the two angles at u, v , neither can exceed π . Therefore the total curvature of the loop does not exceed 2π , and could only achieve 2π when $\kappa_{\text{total}}(\Pi_1) = \kappa_{\text{total}}(\Pi_2) = 0$. Therefore, Π_1 and Π_2 are both straight lines connecting u and v . \square

Lemma 1.6. *Suppose Π_1, Π_2 are two shortest (u, v) -paths in region $R = R[\Pi_1, \Pi_2]$, with $\Pi_1 \cap \Pi_2 = \{u, v\}$. Then each of Π_1, Π_2 touches at least one obstacle in R .*

Proof. Suppose that the conclusion is false. Without loss of generality, Π_1 does not touch any obstacles in R . By Lemma 1.4, the (u, v) -path Π_1 is convex in R . Let Q be the simply connected region obtained by removing the obstacles in R . We have $l(\Pi_1) = l(\Pi_2)$, so they are both shortest (u, v) -paths in the simply connected environment Q . Therefore Π_2 is also convex by Lemma 1.4, if parameterized as a path from v to u . By Lemma 1.5, Π_1 and Π_2 are both straight lines connecting u, v , which contradicts $\Pi_1 \cap \Pi_2 = \{u, v\}$. This proves that when there is more than one shortest (u, v) -path, each of these paths must touch an obstacle. \square

This concludes our topological and geometric preliminaries.

2 Lion's Strategy in a $\text{CAT}(0)$ space

In this section, we describe a winning strategy for a single pursuer in a compact $\text{CAT}(0)$ domain, and prove Theorem 1.1. Our strategy generalizes *lion's strategy* for pursuit in \mathbb{E}^2 , introduced by Sgall [21]. This strategy was adapted for pursuit in polygonal regions by Isler, Kannan and Khanna [14]. Their




Figure 2.1: Lion's strategy in \mathbb{E}^2 . On each move, the pursuer moves on the line segment connecting the center c to the evader, and increases her distance from c .

adaptation relies heavily on the vertices of the polygon P and gives a capture time of $n \cdot \text{diam}(P)^2$, where n is the number of vertices of P . We give a topological version of lion's strategy that succeeds in any compact $\text{CAT}(0)$ domain \mathcal{D} (including polygons) with capture time bounded by $\text{diam}(\mathcal{D})^2$.

Sgall's version of lion's strategy proceeds as follows. Fix a point c as the center of our pursuit, see Figure 2.1. The pursuer starts at $p = c$ and the evader starts at some point e . On her first move, the pursuer moves directly towards e along the line ce . Considering a general round, the pursuer will be on the line segment between c and e prior to the evader move. After the evader moves to $e' \in B(e, 1)$, the pursuer looks at the circle centered at p with radius ϵ . If e is inside this circle, then the pursuer can capture the evader. Otherwise, this circle intersects the line segment ce' at two points. The pursuer moves to the point p' that is closer to e .

Lemma 2.1 (Sgall [21]). *A pursuer using lion's strategy in \mathbb{E}^2 re-establishes her location on the line segment between c and the evader. Furthermore, if the evader moves from e to e' and the pursuer moves from p to p' then $d(c, p')^2 \geq d(c, p)^2 + 1$. \square*

Before generalizing lion's strategy, we introduce the basics of a $\text{CAT}(0)$ geometry; see [10] for a thorough treatment. A complete metric space (X, d) is a *geodesic space* when there is a unique path $\Pi(x, y)$ whose length is the metric distance $d(x, y)$. This path $\Pi(x, y)$ is called a *geodesic* (or shortest path). A *triangle* $\triangle xyz$ between three points $x, y, z \in X$ is the triple of geodesics $\Pi(x, y), \Pi(y, z), \Pi(z, x)$. To each $\triangle xyz \in X$, we associate a *comparison triangle* $\triangle \tilde{x}\tilde{y}\tilde{z} \subset \mathbb{E}^2$ whose side lengths in \mathbb{E}^2 are the same as the lengths of the corresponding geodesics in X . The complete geodesic metric space (X, d) is $\text{CAT}(0)$ when no triangle in X has a geodesic chord that is longer than the corresponding chord in the comparison triangle. In other words, pick any triangle $\triangle xyz$ and any points $u \in \Pi(x, y)$ and $v \in \Pi(y, z)$. Let $\tilde{u} \in \tilde{x}\tilde{y}$ and $\tilde{v} \in \tilde{y}\tilde{z}$ be the corresponding points, chosen distance-wise on the edges $\tilde{x}\tilde{y}$ and $\tilde{y}\tilde{z}$. If the space X is $\text{CAT}(0)$ then $d_X(u, v) \leq d_{\mathbb{E}^2}(\tilde{u}, \tilde{v})$. Colloquially, this is called the “no fat triangles” property, since it also implies that the sum of the angles of the triangle is not greater than π .

Our $\text{CAT}(0)$ lion's strategy generalizes the *extended lion's strategy* for polygons of Isler et al. [14] by taking a more topological approach. The pursuer starts at a fixed center point c and her goal is to stay on the shortest path $\Pi(c, e)$ at all times. In particular, assume that p_t is on the shortest path

$\Pi(c, e_t)$ and that the evader moves from e_t to e_{t+1} . If $d(p_t, e_{t+1}) \leq 1$ then the pursuer responds by capturing the evader. Otherwise, the pursuer draws the unit circle C centered at p_t and moves to the point in $C \cap \Pi(c, e_{t+1})$ that is closest to e_{t+1} .

Lemma 2.2 (Lion's Strategy). *A pursuer using lion's strategy in a $\text{CAT}(0)$ space (X, d) re-establishes her location on the line segment between c and the evader. Furthermore, if the evader moves from e to e' and the pursuer moves from p to p' then $d(c, p')^2 \geq d(c, p)^2 + 1$.*

Proof. Suppose that $p \in \Pi(c, e)$ and then the evader moves to e' . Consider the $\text{CAT}(0)$ triangle $\triangle cee'$ and its comparison triangle $\triangle \tilde{c}\tilde{e}\tilde{e}'$ in \mathbb{E}^2 . Look at the corresponding \mathbb{E}^2 pursuit-evasion game with the pursuer at $\tilde{p} \in \tilde{c}\tilde{e}$. By Lemma 2.1, the pursuer can move to a point $\tilde{p}' \in \tilde{c}\tilde{e}'$ such that $d_{\mathbb{E}^2}(\tilde{c}, \tilde{p}')^2 \geq d_{\mathbb{E}^2}(\tilde{c}, \tilde{p})^2 + 1$. Since there are no fat triangles in X , we have $d_X(p, p') \leq d_{\mathbb{E}^2}(\tilde{p}, \tilde{p}')$ where $p' \in \Pi(c, e')$ is the point corresponding to \tilde{p}' . Therefore, in our original game, the pursuer can move to the point $p' \in \Pi(c, e')$, which satisfies $d_X(c, p')^2 \geq d_X(c, p)^2 + 1$. \square

Finally, we prove Theorem 1.1: lion's strategy succeeds in a $\text{CAT}(0)$ domain.

Proof of Theorem 1.1. Consider a pursuit-evasion game in the compact $\text{CAT}(0)$ domain \mathcal{D} . Pick any $c \in \partial\mathcal{D}$ as our center point. Using lion's strategy, the pursuer increases her distance from c with every step by Lemma 2.2, so she captures the evader after at most $\text{diam}(\mathcal{D})^2$ rounds. \square

3 Minimal Paths and Guarding

The key to our pursuit strategy is the ability of one pursuer to *guard* a shortest path, meaning that the evader cannot cross this path without being caught by a pursuer. When this shortest path splits the domain into two subdomains, the evader will be trapped in a smaller region. We refer to this region as the *evader territory*. In fact, we will be able to also guard a “second shortest path” when the shortest path is already guarded. The definitions and lemmas in this section are adaptations of the minimal path formulations introduced in [8] and further developed in [6]. Recall that we use $d(x, y)$ to denote the length of a shortest (x, y) -path in \mathcal{D} . In addition, we will use $\overset{\circ}{X}$ and \overline{X} to denote the interior and the closure of a set X , respectively.

Definition 2. Let $X \subset \mathcal{D}$ be a path-connected region. The simple path $\Pi \subset X$ is minimal in X when for any $y_1, y_2 \in \Pi$ and any $z \in X$, we have $d_\Pi(y_1, y_2) \leq d_X(y_1, z) + d_X(z, y_2)$.

Definition 3. Let $Z \subset X$ and let Π be a minimal (u, v) -path in Z where $u, v \in \partial Z$. Then the path projection with anchor u is the function $\pi : Z \rightarrow \Pi$ defined as follows. If $d_Z(u, z) < d_Z(u, v) = d_\Pi(u, v)$, then $\pi(z)$ is the unique point $x \in \Pi$ with $d_\Pi(u, x) = d_Z(u, z)$. For the remaining $z \in Z \setminus \Pi$, we set $\pi(z) = v$.

We make a few observations. If $X = \mathcal{D}$, then a shortest (x, y) -path is always a minimal path in \mathcal{D} . In this case, we can define the path projection $\pi : \mathcal{D} \rightarrow \Pi$. When $X \subsetneq \mathcal{D}$, we might have $d_X(x_1, x_2) > d(x_1, x_2)$. In this case, a shortest path in X will be minimal in X , but it will not be minimal in \mathcal{D} . Next, we show that a path projection is non-expansive, meaning that distances cannot increase.

Lemma 3.1. *Let $\pi : Z \rightarrow \Pi$ be a path projection onto a minimal path in Z . Then $d_\Pi(\pi(z_1), \pi(z_2)) \leq d_Z(z_1, z_2)$ for all $z_1, z_2 \in Z$.*

Proof. The proof is a straight-forward argument using the triangle inequality. We consider the case $z_1, z_2 \in Z$ with $d(u, z_1) \leq d(u, z_2) \leq d(u, v)$. We have

$$d_Z(z_2, z_1) \geq d_Z(z_2, u) - d_Z(z_1, u) = d_\Pi(\pi(z_2), u) - d_\Pi(\pi(z_1), u) = d_\Pi(\pi(z_2), \pi(z_1)).$$

The other non-trivial cases are argued similarly. \square

A single pursuer can turn a minimal path Π into an impassable boundary: once the pursuer has attained the position $p = \pi(e)$, the evader cannot cross Π without being captured in response. The proof of the following lemma is similar to the analogous result in [6], but we include this brief argument for completeness.

Lemma 3.2 (Guarding Lemma). *Let $\pi : X \rightarrow \Pi$ be a path projection onto the minimal (u, v) -path $\Pi \subset X$. Consider a pursuit-evasion game between pursuer p and evader e in the environment X .*

- (a) *After $O(\text{diam}(X))$ turns, the pursuer can attain $p^t = \pi(e^{t-1})$.*
- (b) *Thereafter, the pursuer can re-establish $p^{s+1} = \pi(e^s)$ for all $s \geq t$.*
- (c) *If the evader moves so that a shortest path from e^{s-1} to e^s intersects Π , then the pursuer can capture the evader at time $s + 1$.*

Proof. To achieve (a), the pursuer moves as follows. First, p travels to u , reaching this point in $O(\text{diam}(X))$ turns. Next, p traverses along Π until first achieving $d(u, p^i) \leq d(u, \pi(e^{i-1})) < d(u, p^i) + 1$. If $p^i = \pi(e^{i-1})$ then we are done. Otherwise, when the evader moves, we either have $d(u, p^i) - 1 < d(u, \pi(e^i)) \leq d(u, p^i) + 1$ or $d(u, p^i) + 1 < d(u, \pi(e^i)) < d(u, p^i) + 2$ by Lemma 3.1. In the former case, p can move to $\pi(e^i)$ in response, achieving her goal. In the latter case, p will increase her distance from u by one unit, re-establishing $d(u, p^{i+1}) \leq d(u, \pi(e^i)) < d(u, p^{i+1}) + 1$. This latter evader move can only be made $O(\text{diam}(X))$ times, after which the pursuer achieves $p = \pi(e)$.

Next, suppose that $p^s = \pi(e^{s-1})$ and that $e^{s-1} \in X \setminus \Pi$. The pursuer can stay on the evader projection by induction since

$$d_\Pi(p^t, \pi(e^t)) = d_\Pi(\pi(e^{t-1}), \pi(e^t)) \leq d(e^{t-1}, e^t) \leq 1,$$

so (b) holds. As for (c), suppose that a shortest path from e^{s-1} to e^s includes the point $y \in \Pi$. Then

$$d(p^t, e^t) \leq d_\Pi(\pi(e^{t-1}), y) + d(y, e^t) \leq d(e^{t-1}, y) + d(y, e^t) = d(e^{t-1}, e^t) = 1.$$

Therefore the pursuer can capture the evader on her next move. \square

The Guarding Lemma is the cornerstone of our pursuer strategy. When the pursuer moves as specified in the lemma, we say that she *guards* the path Π . In Section 4, our pursuers will repeatedly guard paths chosen to reduce the number of obstacles in the evader territory. Once the evader is trapped in a region that is obstacle-free, we have reached the endgame of the pursuit.

Lemma 3.3. *Suppose that the evader is located in a simply connected region R whose boundary consists of subcurves of the original boundary $\partial\mathcal{D}$ and two guarded paths Π_1 and Π_2 . If the evader remains in R , then the third pursuer can capture him in finite time. If the evader tries to leave the region through Π_1 or Π_2 , then he will be captured by the guarding pursuer.*

Proof. By Lemma 3.2, if the evader tries to leave this region, he will be caught by either p_1 or p_2 . If the evader remains in this component, then Theorem 1.1 guarantees that pursuer p_3 captures the evader in a finite number of moves. \square

The remainder of this section is devoted to identifying guardable paths that touch obstacles. Guarding such a path will neutralize the threat posed by the obstacle. First, we consider the case when p_1 guards the unique shortest (u, v) -path Π that touches an obstacle O in the evader region. The objective of p_2 is to guard another (u, v) -path Π_2 of a different homotopy type. This path can be guarded even when Π_2 is longer than Π_1 , provided that any path shorter than Π_2 also intersects Π_1 .

Lemma 3.4. *Suppose that the evader territory $R = [\Pi_1, \Delta]$ is bounded by the unique (u, v) -shortest path Π_1 and another boundary curve Δ . Furthermore, suppose that Π_1 touches an obstacle O and that Π_1 is guarded by p_1 . Then we can find a (u, v) -path $\Pi_2 \subset R$ with the following properties: (a) $O \subset R[\Pi_1, \Pi_2]$, so that the homotopy type of Π_2 is different than that of Π_1 ; and (b) Π_2 is guardable by p_2 , provided that Π_1 remains guarded by p_1 .*

A naive attempt to find such a path is to pick some $x \in \Pi_1 \cap \bar{O}$ and find a shortest path that does not include the point x . However, $R \setminus \{x\}$ is not a closed set, which would complicate our argument. Furthermore, it could be that the next shortest path includes x without using this point as a shortcut around the obstacle O , as shown in Figure 3.1(b).¹ We handle both problems by removing a small and well-chosen open region A near x , rather than removing the point x . The delicate choice of A relies on the existence of both κ_{\max} and d_{\min} .

¹We note that this unusual circumstance is overlooked in [8], where it can occur during their minimal path strategy. This case can be easily handled in a manner analogous to our approach, but based on the visibility graph of their environment.




Figure 3.1: Finding the second shortest path. (a) When $\Pi \cap \partial O$ contains a curve, we can remove a small open ball. (b) The shortest path Π touches obstacle O at x . The second shortest path Π goes around O , but includes the point x . (c) Finding Π_2 requires removing a small, open, triangular set A between Π_1 and O , and then finding the shortest (u, v) -path in the closed set $R \setminus A$. (d) Any path that crosses the line segment yz can be short-cut.

Proof. First, suppose that $\Pi_1 \cap \partial O$ includes a continuous subcurve $C \subset \Pi_1$. Pick $x \in \mathring{C}$ and $\epsilon > 0$ so that $B(x, \epsilon) \cap R \subset \bar{O}$. Let $R' = R \setminus \mathring{B}(x, \epsilon)$, which effectively absorbs the obstacle into the boundary, see Figure 3.1 (a). The region R' is closed, so there is a well-defined shortest (u, v) -path $\Pi_2 \subset R'$. The path Π_2 is guardable in R' , and therefore it is guardable by p_2 in R , provided that p_1 guards Π_1 . Indeed, any shorter path in R must go through the point x , so Lemma 3.2 guarantees that an evader using such a path will be caught by p_1 . Finally, we note that Π_1, Π_2 have distinct homotopy types because $O \subset R[\Pi_1, \Pi_2]$.

Next, we consider the case where $\Pi_1 \cap \bar{O}$ contains no continuous curves: we just focus on the first point $x \in \Pi_1 \cap \bar{O}$ encountered as we move from u to v . Locally around x , the path Π_1 and the boundary ∂O separate R into two external regions (outside of Π and inside ∂O) and two internal regions, see Figure 3.1 (b). The shortest path Π_1 does not self-intersect, so locally near x , this path consists of two line segments meeting at x , create an interior angle smaller than 2π . Therefore, at least one of the two interior angles made by Π_1 and the obstacle tangent line(s) at x is strictly less than π . This local region is where we will remove our triangular open set.

Without loss of generality, suppose that the subpath $\Pi_1(x, v)$ helps to bound this local region. Take points $y \in \Pi_1(x, v)$ and $z \in \partial O$ (traveling counterclockwise from x) such that $0 < d_{\Pi_1}(x, y) = d_{\partial O}(x, z) < d_{\min}/2$, and the angle $\angle yxz < \pi$. Let $A \subset B(x, d_{\min})$ be the closed region with endpoints (x, y, z) , where the third curve is the unique shortest (y, z) -path Γ , see Figure 3.1 (c). The bound on the absolute curvature κ_{\max} allows us to choose our y, z so that the region A is essentially triangular. Since d_{\min} is the minimum distance between obstacles, A is obstacle-free, so Γ is a straight line segment.

We remove the relatively open set $A' = A \setminus \Gamma$ from our domain. We then find the shortest (u, v) -path Π_2 in the closed set $R = R \setminus A'$. We claim that p_2 can guard Π_2 in R , provided that p_1 guards Π_1 . As in the previous case, the shorter paths that go through x are not available to the evader. Therefore, we must show that any path in R that visits A' is longer than Π_2 . Such a path Λ must enter and leave A' through Γ , say at points a, b , see Figure 3.1 (d). However, the subpath $\Lambda(a, b)$ can be replaced with the unique shortest path $\Gamma(a, b)$ without changing the homotopy type, a contradiction. Once again, Π_1, Π_2 have distinct homotopy types because $O \subset R[\Pi_1, \Pi_2]$. \square

We refer to the paths Π_1, Π_2 from Lemma 3.4 as a *guardable pair*. Provided that the shortest (u, v) -path Π_1 is guarded, the “second shortest (u, v) -path” Π_2 can also be guarded. The following corollary is a variation of the lemma.

Corollary 3.5. *Let Π_1, Π_2 be (u, v) -paths that are guarded by p_1, p_2 , respectively. Suppose that for $i = 1, 2$, the path Π_i touches an obstacle O_i , where $O_1 \neq O_2$. Then we can find a path Π_3 with the following properties: (a) the homotopy type of Π_3 is different than the homotopy types of Π_1 and Π_2 , and in particular, $O_i \in R[\Pi_i, \Pi_3]$ for $i = 1, 2$; and (b) Π_3 is guardable by p_3 , provided that Π_1, Π_2 remain guarded by p_1, p_2 .*

Proof. The proof is similar to the proof of Lemma 3.4. This time, we must remove an open set A near $x \in \Pi_1 \cap O_1$ and an open set B' near $y \in \Pi_2 \cap O_2$. We then find Π_3 in $R \setminus (A' \cup B')$. \square

This concludes our search for guardable paths that touch obstacles. The next section lays out the three-pursuer strategy for capturing the evader in a two-dimensional domain.

4 Shortest Path Strategy: Proof of Theorem 1.2

In this section, we prove Theorem 1.2: three pursuers can capture an evader in a two-dimensional compact domain with piecewise analytic boundary. We adapt the the shortest path strategy of Bhaudaria et al. [8] to our more general topological setting. In particular, our guardable path lemmas from Section 3 supplant their use of polygon vertices to find successive paths. Their algorithm guarantees success by reducing the number of polygon vertices in the evader territory. Instead, we keep track of the threat level of obstacles to argue that the evader becomes trapped in a simply connected region.

Our pursuit proceeds in rounds. At the start of a round, at most two pursuers guard paths. The third pursuer moves to guard another path with the goal of eliminating obstacles from the evader territory. This third path will either be a shortest path, or it will create a guardable pair with the currently guarded path(s). Once this third path is guarded, the evader is trapped in a smaller region, which releases one of the other pursuers to continue the process. This continues until the evader is trapped in a simply connected region, where the free pursuer can capture the evader by Lemma 3.3.

We start by showing that the boundary of the evader territory is always piecewise analytic, after recalling two definitions. First, the endpoints of a line segment touching the boundary $\partial\mathcal{D}$ are called *switch points*. Second, a point x is an *accumulation point* (or limit point) of a set S when any open set containing x contains an infinite number of elements in S . We make use of the following result about the interaction of a geodesic with the boundary of the domain.

Theorem 4.1 (Albrecht and Berg [2]). *If M is a 2-dimensional analytic manifold with boundary embedded in \mathbb{E}^2 , and γ is a geodesic in M , then the switch points on γ have no accumulation points.*

We restrict ourselves to analytic boundary, instead of smooth (C^2 , or even C^∞) boundary, to avoid some potentially pathological behavior of geodesics. For example, Albrecht and Berg [2] construct a geodesic in C^∞ environment, that achieves a Cantor set of positive measure as the accumulation of switch points. This unusual geometry hampers our ability to confine the evader in a well-defined connected component. Theorem 4.1 ensures that our new evader territory will be bounded by piecewise analytic curves.

Lemma 4.2. *Let \mathcal{D} be a compact domain with piecewise analytic boundary. If Π is a shortest path in \mathcal{D} , then Π is piecewise analytic. Furthermore, if $\mathcal{D} \setminus \Pi$ is disconnected, then it contains finitely many connected components, and the boundary of each connected component is piecewise analytic.*

Proof. Let $B \subset \partial(\Pi \setminus \partial\mathcal{D})$ be the set of *switch points*. We claim that B is finite. Otherwise, there must be an accumulation point of B since $l(\Pi)$ is finite, contradicting Theorem 4.1. Now we can use the finite set B as endpoints to partition Π so that each subcurve is either in the boundary $\partial\mathcal{D}$, or in the interior $\mathring{\mathcal{D}}$. Since any shortest path in the interior $\mathring{\mathcal{D}}$ must be a line segment, the path Π is piecewise analytic. For each connected component of $\mathcal{D} \setminus \Pi$, its boundary is a subset of $\Pi \cup \partial\mathcal{D}$, hence is piecewise analytic. \square

In order to prove Theorem 1.2, we will show that our pursuit succeeds in finite time. To aid in this effort, we assign a threat level to each of the k obstacles in the original domain. These threat levels will reliably decrease during pursuit. An obstacle is in one of three states: dangerous, safe, or removed. A *removed* obstacle lies outside the evader territory. A *safe* obstacle lies in the evader territory and touches a currently guarded path. This obstacle is not a threat because the evader cannot circle around the object without being captured. The remaining obstacles are *dangerous*. Finally, we say that the evader territory is dangerous if it contains at least one dangerous obstacle.

At the start of pursuit, all obstacles are dangerous. So long as there are still dangerous obstacles, a round consists of taking control of a guardable path. This effort succeeds in a finite number of moves by Lemma 3.2. We will show that after at most two rounds, either a dangerous obstacle transitions to safe/removed, or a safe obstacle transitions to removed. This is our notion of progress: after at most $2k$ rounds, the evader territory is not dangerous. From here forward, we focus on the transition of the threat levels of obstacles.




Figure 4.1: The shortest (u, v) -path Π_1 guarded in round one. (a) The path Π_1 partitions the outer boundary to subcurves Δ_1, Δ_2 of equal length. (b) The evader may be trapped in a pocket between the path Π_1 and the boundary subcurve Δ_3 of obstacle O .

In general, our evader territory will be bounded by part of the domain boundary $\partial\mathcal{D}$ and by at most two guarded paths Π_1, Π_2 . At the end of a round, the evader territory will be updated, bounded in part by updated paths Π'_1, Π'_2 . If these guarded paths intersect or share subpaths, then the evader is actually trapped in a smaller region by Lemma 3.2. When this is the case, we advance the endpoint(s) of our paths so that these are the only point(s) shared by our paths. This obviates the need to discuss degenerate cases.

The first round is an initialization round, so all obstacles might still be dangerous when this round completes. However, we will be able to neutralize at least one obstacle in the subsequent round. To kick off the first round, we pick points $u, v \in \partial\mathcal{D}$, chosen so that they divide the outer boundary into two curves Δ_1, Δ_2 of equal length, see Figure 4.1(a). Let Π_1 be a shortest (u, v) -path; if there are multiple shortest paths (in which case each touches an obstacle), then we pick one arbitrarily. Using Lemma 3.2, p_1 moves to guard Π_1 . The round ends when p_1 has attained guarding position, trapping the evader in a subdomain that is bounded by Π_1 and one of Δ_1, Δ_2 . The evader could be trapped in a smaller pocket region between Π_1 and a subcurve Δ_3 of an obstacle $O \subset \mathcal{D}$, see Figure 4.1(b). In the latter case, the obstacle O is marked as removed and we treat Δ_3 as the new outer boundary. After updating the evader territory R , any obstacle $O \not\subset R$ is marked as removed. Any obstacle $O \subset R$ that touches Π_1 or Π_2 is marked as safe.

For the remainder of the game, the boundary of the evader territory is one of the following types.

- Type 0 region: A region containing no dangerous obstacles.
- Type 1 region: A dangerous three-sided region bounded by a (u, v) -shortest path Π_1 , a (u, w) -shortest path Π_2 and a (v, w) -path $\Delta \subset \partial\mathcal{D}$. No obstacle touches both Π_1, Π_2 .
- Type 1' region: A dangerous two-sided region bounded by a (u, v) -shortest path Π_1 and a (u, v) -path $\Delta \subset \partial\mathcal{D}$. We treat this as a special case of the previous type, where Π_2 consists of the single point $w = u$. This point is on Π_1 , so it is guarded by p_1 .

- Type 2 region: A dangerous two-sided region bounded by (u, v) -paths Π_1, Π_2 , each of which touches an obstacle in the evader territory. The path Π_1 is a shortest (u, v) -path in this region. The path Π_2 might also be a shortest (u, v) -path, or it could be a “second shortest path,” meaning that it is a shortest (u, v) -path among the set of (u, v) -paths that are not homotopic to Π_1 . No obstacle touches both Π_1, Π_2 .
- Type 3 region: a dangerous 4-sided region bounded by a (u, v) -shortest path Π_1 , a (w, x) -shortest path Π_2 , a (v, w) -path Δ_1 from the boundary and a (u, x) -path Δ_2 from the boundary. These vertices are arranged so that they are ordered clockwise as u, v, w, x . No obstacle touches both Π_1, Π_2 .

For example, after the initialization round, the evader territory is a type 1' region, bounded by a guarded path and part of the boundary $\partial\mathcal{D}$. Finally, we emphasize that Lemma 4.2 ensures that the boundary of the evader region is always piecewise analytic, since it consists of sub-curves of the piecewise analytic boundary along with one or more shortest paths.

We now describe the different types of rounds. In regions of type 1, 1' and 2, we will always transition at least one obstacle. At the end of such a round, the evader could now be trapped in a region of any type. Type 3 rounds are slightly different. Our primary goal is to trap the evader in a type 1 region, where we will surely make progress in the subsequent next round. However, it is possible to transition an object via a type 3 move (just as in the initialization round). In this case, we make immediate progress, and the evader could then be trapped in a region of any type.

First we consider type 1 regions. This also handles type 1' regions as a special case. We use the following lemma to identify a point $x \in \Delta$ and a shortest (u, x) -path to guard during this round.

Lemma 4.3. *Let shortest paths $\Pi_0(u, v), \Pi_1(u, w)$ and boundary path $\Delta(v, w)$ bound a type 1 environment R . If R contains obstacles, then there exists a point $x \in \Delta$ such that there are multiple shortest (u, x) -paths in R , each of which touches at least one obstacle.*

Proof. Parameterize the boundary path as $\Delta : [0, 1] \rightarrow R$. We prove the lemma by contradiction.

Suppose that for every $t \in [0, 1]$, the shortest $(u, \Delta(t))$ -path Π_t is unique. Denote its length by $l(\Pi_t) = d(u, \Delta(t))$. Define the function $n(t)$ to be the number of obstacles in the region R_t bounded by Π_0, Π_t and Δ . The function $n(t)$ is well-defined by the uniqueness of each Π_t . Furthermore, $n(0) = 0$, and $n(1) > 0$, so there must be a jump discontinuity somewhere in $[0, 1]$. Let $s = \inf\{t \in [0, 1] : n(t) > 0\}$.

Case 1: $n(s) > 0$ where $0 < s \leq 1$. (Recall that $n(0) = 0$, so $s > 0$.) Let Γ be a shortest $(u, \Delta(s))$ -path that is of the same homotopy class as Π_0 . The choice of s and the uniqueness of Π_s guarantee that $l(\Gamma) > l(\Pi_s)$. Also no obstacles are contained in the region bounded by Π_0, Γ and $\Delta(0, s)$. By the continuity of $l(\Gamma_t) = d(u, \Delta(t))$ with respect to t , we have $l(\Gamma) \leq l(\Pi_t) + l(\Delta(t, s))$,




Figure 4.2: Representative examples of a type 1 move, where we transition to (a) a type 1 region, (b) a type 1 or type 1' region, (c) a type 1 or a type 3 region, (d) a type 1 or type 2 region, (e) a type 1 or type 3 region.

so

$$l(\Gamma) \leq \lim_{t \rightarrow s^-} (l(\Pi_t) + l(\Delta(t, s))) = l(\Pi_s) + 0 = l(\Pi_s).$$

This contradicts $l(\Gamma) > l(\Pi_s)$, so Π_s is not the unique shortest $(u, \Delta(s))$ -path.

Case 2: $n(s) = 0$, where $0 \leq s < 1$. Let $\{s_i\}$ be an infinite sequence $s_i \rightarrow s^+$, such that $n(s_i) > 0$ for all i . There are finite number of obstacles, so by taking a subsequence if necessary, we can assume that the shortest paths $\{\Pi_{s_i}\}$ are of the same homotopy class. Let Γ be the shortest $(u, \Delta(s))$ -path of this homotopy class. We have $l(\Gamma) \leq l(\Pi_{s_i}) + l(\Delta(s, s_i))$ for all i , and therefore

$$l(\Gamma) \leq \lim_{i \rightarrow \infty} (l(\Pi_{s_i}) + l(\Delta(s, s_i))) = l(\Pi_s),$$

where the limit holds by the continuity of distances in the region. However, this contradicts the uniqueness of Π_s which would require $l(\Pi_s) < l(\Gamma)$.

Finally, we can conclude that there are multiple shortest (u, x) -paths. By Lemma 1.6, each of these paths touches at least one obstacle. \square

Having found the next path to guard, we now prove that we transition an object during a type 1 move.

Lemma 4.4. *Suppose that the evader is trapped in a type 1 (or type 1') region. Then the third pursuer can guard a path that transitions an obstacle state.*

Proof. By Lemma 4.3, there is some point $x \in \Delta$ with multiple shortest (u, x) -paths, each of which touches an obstacle. Let Π_3 be one of these shortest (u, x) -paths. If $x = v$ then we take a path $\Pi_3 \neq \Pi_1$. Similarly if $x = w$ we choose $\Pi_3 \neq \Pi_2$. When $x \notin \{v, w\}$, we can choose Π_3 arbitrarily from the collection of (u, x) -shortest paths. Pursuer p_3 moves to guard Π_3 , which traps the evader in either $R[\Pi_1, \Pi_3]$ or $R[\Pi_3, \Pi_2]$. Any obstacles in the other region are marked as removed.

Without loss of generality, let $O \subset R[\Pi_1, \Pi_3]$ be an obstacle touched by Π_3 , see Figure 4.2(a). Suppose that prior to p_3 guarding Π_3 , the object O was dangerous. If $e \in R[\Pi_3, \Pi_2]$ then O transitions to removed. If $e \in R[\Pi_1, \Pi_3]$ then O transitions to safe. However, we may be in a more advantageous position, shown in Figure 4.2(b): the evader could be trapped in a pocket between obstacle O and path Π_3 . In this case, the new evader territory is type 1' and the obstacle O is marked as removed, since it is now part of the outer boundary of the evader territory.

Next, suppose that O was already safe, touched by Π_1 . If $e \in R[\Pi_2, \Pi_3]$ after p_3 guards Π_3 , then O transitions from safe to removed. If the evader is trapped in a pocket region between O and Π_3 , we proceed as in the previous case. Otherwise, we have $e \in R[\Pi_1, \Pi_3]$ and the obstacle O separates $R[\Pi_1, \Pi_3]$ into disjoint regions, as shown in Figure 4.2(c). The evader is trapped in one of these two subregions because both Π_1, Π_3 are guarded. Let Δ' be the subcurve of ∂O that bounds the effective evader territory. We update the evader territory appropriately, bounded by Δ' and subpaths of Π_1, Π_3 , and perhaps part of Δ . The result is a region of type 1 or 3. The obstacle O is marked as removed: it is now part of the boundary. This reduces the number of safe obstacles.

When Π_3 touches multiple obstacles, each of them transitions to a lower threat level. Figures 4.2(d) and (e) show that we can also end up in a type 2 or 3 region, depending on the configuration of these obstacles and the location of the evader at the end of the round. \square

Next, we consider a type 2 region. Such a region is bounded by (u, v) -paths Π_1, Π_2 that form a guardable pair, where Π_1, Π_2 touch safe obstacles O_1, O_2 , respectively. Without loss of generality, Π_1 is a shortest (u, v) -path in the region, and Π_2 is either another shortest path, or a “second shortest path” as found in Lemma 3.4. (A type 1 move can lead to the first case. A type 2 move can lead to the second case, as we are about to see.)

Lemma 4.5. *Suppose that the evader is trapped in a type 2 region. Then the third pursuer can guard a path that transitions an obstacle state.*

Proof. Use Corollary 3.5 to find a guardable (u, v) -path Π_3 in $R[\Pi_1, \Pi_2]$ whose homotopy type is distinct from that of both Π_1, Π_2 . Pursuer p_3 establishes a guarding position on Π_3 . The evader is now trapped in either $R[\Pi_1, \Pi_3]$ or $R[\Pi_3, \Pi_2]$, so one of O_1, O_2 transitions from safe to removed. Furthermore, Π_3 must touch at least one obstacle in each of $R[\Pi_1, \Pi_3]$ or $R[\Pi_3, \Pi_2]$. Otherwise, Π_3 would be shorter than one of Π_1, Π_2 , which contradicts the minimality of that path in $R[\Pi_1, \Pi_2]$. Depending on the configuration of the obstacles, we may be able to restrict the evader territory further. After doing so, the evader territory may be of any possible type, as shown in Figure 4.3 (a). \square

This brings our discussion to a type 3 region, with p_1 guarding shortest (u, v) -path Π_1 and p_2 guarding shortest (w, x) -path Π_2 . Our primary goal is to trap the evader in a type 1 region, but we might end up transitioning an obstacle instead. In the latter case, the new evader territory can be of any type, as explained below.




Figure 4.3: Examples of moves where the new guarded path Π_3 divides the region into five subregions, each identified by its type. (a) A type 2 move. (b) A type 3 move.

Lemma 4.6. *Suppose that the evader is trapped in a type 3 region. Then the third pursuer can guard a path so that either (a) the evader is trapped in a type 1 region, or (b) an obstacle transitions to a lower threat level.*

Proof. Let Π_3 be a (u, w) -shortest path. Pursuer p_3 moves to guard this path using Lemma 3.2. This traps the evader in a smaller region: without loss of generality, this region is bounded by Π_1, Δ_1, Π_3 . If Π_3 does not touch any obstacles in this region, then the evader is now in a type 1 region. If Π_3 touches an obstacle O , then this obstacle transitions to either safe or removed. The evader could be trapped in a region of any type, as shown in Figure 4.3 (b). \square

We can now prove our main theorem: three pursuers can capture the evader in a two-dimensional compact domain with piecewise analytic boundary.

Proof of Theorem 1.2. The first round traps the evader in a type $1'$ region, or transitions an obstacle state. If we are in a region of type 1, $1'$ or 2 then we transition an obstacle state in the current round by Lemma 4.4 and Lemma 4.5. When we are in a type 3 region, Lemma 4.6 ensures that we either trap the evader in a type 1 region, or we transition an obstacle. With each path that we guard, the boundary of the updated evader territory is still piecewise analytic by Lemma 4.2. At the end of the round, we update the evader territory and our value for minimum obstacle separation since our new guarded path might be closer to an obstacle than the current value d_{\min} . (Note that the maximum boundary curvature κ_{\max} never increases since all additions to the boundary are line segments.)

After at most $2k$ rounds, we have transitioned all k obstacles to either safe or removed. Once all obstacles have been transitioned, the evader is trapped in a simply connected type 0 region. Lemma 3.3 shows that the evader will then be caught. Each round completes in finite time, so the three pursuers win the game. The capture time upper bound of $O(2k \cdot \text{diam}(\mathcal{D}) + \text{diam}(\mathcal{D})^2)$ follows easily. The time required to guard any shortest path is $\text{diam}(\mathcal{D})$ by Lemma 3.2 and lion's strategy completes in time $\text{diam}(\mathcal{D})^2$ by Theorem 1.1. \square

5 Conclusion

In this paper, we described a winning pursuer strategy for a single pursuer in a $\text{CAT}(0)$ space when the capture criterion is $d(p, e) = 0$. We then restricted our attention to compact domains in \mathbb{E}^2 with piecewise analytic boundary. We showed that three pursuers are sufficient to catch an evader in such environments.

There are plenty of avenues for research in topological pursuit-evasion. Pursuit-evasion results on polyhedral surfaces are an active area of current research [15, 18, 19]. For example, Klein and Suri [15] have proven that $4g + 4$ pursuers have a winning strategy on a polyhedral surface of genus g . Meanwhile, Schröder [20] has proven that at most $\lfloor 3g/2 \rfloor + 2$ pursuers are needed for a graph of genus g (meaning that such a graph can be drawn on a surface of genus g without edge crossings). It would be natural to consider this question for topological surfaces. Likewise, there are a wealth of motion and sensory constraints to consider. Most of these variations of pursuit-evasion have a natural analog in a topological setting.

6 Acknowledgments

This work was completed while the first author was a long-term visitor at the Institute for Mathematics and its Applications. We are grateful to the IMA for its support and for fostering such a vibrant research community.

References

- [1] M. Aigner and M. Fromme. A game of cops and robbers. *Discrete Appl. Math.*, 8:1–12, 1983.
- [2] F. Albrecht and I. D. Berg. Geodesics in euclidean space with analytic obstacles. *Proceedings of the American Mathematical Society*, 113(1):201–207, 1991.
- [3] S. Alexander, R. Bishop, and R. Ghrist. Pursuit and evasion in non-convex domains of arbitrary dimensions. In *Proceedings of Robotics: Science and Systems*, 2006.
- [4] S. Alexander, R. Bishop, and R. Ghrist. Capture Pursuit Games on Unbounded Domains. *L'Enseignement Mathématique*, 55:103–125, 2009.
- [5] S. Alexander, R. Bishop, and R. Ghrist. Total Curvature and Simple Pursuit on Domains of Curvature Bounded Above. *Geometriae Dedicata*, 149(1):275–290, 2010.
- [6] B. Ames, A. Beveridge, R. Carlson, C. Djang, V. Isler, S. Ragain, and M. Savage. A leapfrog strategy for pursuit-evasion in polygonal environments. *International Journal of Computational Geometry and Applications*, to appear.

- [7] M. Armstrong. *Basic Topology*. Springer, 1983.
- [8] D. Bhadauria, K. Klein, V. Isler, and S. Suri. Capturing an evader in polygonal environments with obstacles: The full visibility case. *International Journal of Robotics Research*, 31(10):1176–1189, 2012.
- [9] A. Bonato and R. Nowakowski. *The Game of Cops and Robbers on Graphs*. American Mathematical Society, 2011.
- [10] M. Bridson and A. Haefliger. *Metric Spaces of Non-positive Curvature*. Springer-Verlag, 1999.
- [11] T. H. Chung, G. A. Hollinger, and V. Isler. Search and pursuit-evasion in mobile robotics: a survey. *Autonomous Robots*, 31:299–316, 2011.
- [12] M. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.
- [13] A. Hatcher. *Algebraic Topology*. Cambridge Univeristy Press, 2002.
- [14] V. Isler, S. Kannan, and S. Khanna. Randomized pursuit-evasion in a polygonal environment. *IEEE Transactions on Robotics*, 5(21):864–875, 2005.
- [15] K. Klein and S. Suri. Pursuit-evasion on polyhedral surfaces. In L. Cai, S. W. Cheng, and T.-W. Lam, editors, *Algorithms and Computation: 24th International Symposium*, pages 284–294. Springer, 2013.
- [16] S. Kopparty and C. V. Ravishankar. A framework for pursuit evasion games in \mathbb{R}^n . *Inf. Process. Lett.*, 96:114–122, 2005.
- [17] J. E. Littlewood. *Littlewood’s Miscellany*. Cambridge Univeristy Press, 1986.
- [18] N. Noori and V. Isler. Lion and man game on convex terrains. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2014.
- [19] N. Noori and V. Isler. Lion and man game on polyhedral surfaces with boundary. In *IEEE Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [20] B. S. Schröder. The copnumber of a graph is bounded by $\lfloor \frac{3}{2}\text{genus}(G) \rfloor + 3$. In J. Koslowski and A. Melton, editors, *Categorical Perspectives: Proceedings of the Conference in Honor of George Strecker’s 60th Birthday*, pages 243–263, 2001.
- [21] J. Sgall. A solution to David Gale’s lion and man problem. *Theoretical Comp. Sci.*, 259(1–2):663–670, 2001.